When Willow Blooms: Category-Theoretic and Homotopic Pathways Toward the Collapse of Cryptographic Resistance

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 23, 2025

Google's Willow processor just demonstrated a *verifiable* quantum advantage using a new "Quantum Echoes" algorithm that measures out-of-time-order correlators (OTOCs) with a forward–perturb–reverse protocol. The task was verified by echoing the computation (run U, perturb, then run U^{-1}) so that correctness emerges as a physically checkable "return signal," and the team reports a ≈13,000× speedup over leading classical approaches for this class of correlators. This matters because *verification* converts quantum outputs from curiosities into artifacts one can trust and reuse—for molecular modeling, Hamiltonian learning (e.g., NMR), and other physics workloads. We place Willow's result into a categorical/homotopic lens. Classical cryptographic hardness—e.g., SHA-256 preimages—is a claim about *shape*: one-way morphisms without sections in a topos whose fibrations encode feasible computation. Quantum Echoes does not "iterate faster" so much as *retype the space*: echo verifiability acts like a functorial lifting that creates computable witness paths where none existed classically. While Willow doesn't break modern cryptography today (qubit counts and fault tolerance still bound practical attacks), the structural move—verifiable, physically grounded witness construction—signals how quantum structure can flatten hard types without targeting a specific scheme. Post-quantum security, then, should be cast as *topos-invariant*: security that survives adversarial retypings, not merely larger parameters.

Keywords: quantum echoes, OTOCs, Willow chip, verifiable quantum advantage, SHA-256, Grover's algorithm, category theory, homotopy type theory, topos theory, sheaves, fibrations, witness paths, Hamiltonian learning, NMR, post-quantum cryptography.

A collaboration with GPT-4o and GPT-5. CC4.0.

Outline

1. Reframing Cryptographic Resistance as Shape
    - From cost to *cohomology*: resistance as obstruction classes in a classical topos.
    - One-way functions as non-split epis; preimage search as section non-existence.
    - Why *verifiable* quantum tasks change the epistemic status of outputs.
2. Category Theory of One-Way Functions
    - Faithful vs. full functors: compression morphisms and loss of invertible data.
    - SHA-256 as a collapse morphism (entropy reduction) and fibrational boundaries.
    - Obstruction to global sections as the categorical face of "hardness."
3. HoTT of the Preimage Problem
    - Dependent types for instances–witnesses; uninhabited preimage types classically.
    - "Constructive existence" vs. mere existence: when a witness is physically realizable.
    - Amplitude amplification as curvature on a type space (intuition for Grover).
4. Grover's Algorithm as Functorial Lifting
    - Grover as a diagonal on subobjects; square-root law as height collapse.
    - From search cost to path deformation—homotopies that shorten certificate paths.
    - Limits: oracle assumptions; why Grover alone doesn't trivialize hashes.
5. Willow and Quantum Echoes: Echo-Verifiable Witness Construction
    - The Echo protocol: run U, perturb, run U^{-1}; compare to measure OTOCs.
    - Why echoing enables *verifiability* (Loschmidt-style return as a physical section).
    - Empirics: ≈13,000× speedup for OTOC(2)-like tasks; Nature publication claims.
    - Practical hooks: Hamiltonian learning/NMR and molecular structure inference.
6. Echoes as Zigzag Diagrams in the Quantum Topos
    - The $F \circ P \circ F^{-1}$ loop as a zigzag witnessing equivalence classes.
    - Noise as a sheaf deformation; logical qubits as near-global sections under refinement.
    - Echo-stability criteria as homotopy invariants of computation.

7. Collapse Conditions for Cryptographic Types
    - When obstruction classes vanish: from hard preimage types to contractible ones.
    - Distinguish “break by qubits” from “retype by structure”: echoes vs. oracles.
    - Adversarial adjunctions that turn non-sections into sections (threat model upgrade).
8. Post-Quantum Security as Topos-Invariance
    - Protocols as objects in a fibred category; security as preservation under retyping.
    - Sheaf-theoretic views of commitment, randomness, and transcript consistency.
    - Design goal: primitives whose hardness is stable under echo-like liftings.
    - Standards context and near-term reality: Willow doesn’t break RSA today; PQC remains prudent.
9. From Resonance to Echoes: Verification as Epistemic Upgrade
    - Why *verifiable* advantage matters more than raw advantage for downstream science.
    - Data you can *trust*: enabling virtuous loops in AI, chemistry, and materials.
10. The Cryptographic Singularity, Recast

- Not a date but a limit object: as echo-style functors proliferate, hard types flatten.
- Predictive markers: growth of echo-verifiable workloads; cross-domain witnesses.

11. Conclusion: Secure the Shape, Not the Speed

- Security as geometric robustness to adversarial retypings.
- Willow as an early witness that *verification* is the decisive axis.
- Program of work: topos-stable primitives; echo-aware threat models; composable proofs.

References

Reframing Cryptographic Resistance as Shape

From cost to cohomology: resistance as obstruction classes in a classical topos

Treat a cryptosystem as a diagram in a classical "computational topos" T whose objects are feasible information types and whose morphisms are computable maps under a fixed physical model (classical circuits, RAM, etc.). A hash h: X -> Y is a compression morphism that forgets microstructure. In T, *security* is not merely large expected work; it is the *non-existence of certain lifts and sections*, captured abstractly as obstruction classes.

Concretely, let p: E -> X encode feasible preimage families for a public output y in Y, with a section s: X -> E corresponding to a constructive witness picker. The obstruction to a global section lives in a cohomology class [o] in H^1(T, F) for an appropriate sheaf F of local witnesses: local trivialisations (short proofs, collisions in small neighborhoods, partial inverses under constraints) may exist, yet fail to glue globally. "Hardness" then reads as persistence of [o] under all classically allowed refinements: even after refining covers (more memory, more time, better heuristics), the class does not vanish. Thus resistance is geometric: a stable obstruction to gluing local evidence into a global, computable witness.

This lens matters now because *verifiable* quantum protocols change the sheaf of evidence itself (what counts as a locally checkable witness), potentially altering the obstruction landscape without directly "breaking" a particular scheme. Google's *Quantum Echoes* result is precisely about turning certain quantum outputs into physically verifiable evidence, which upgrades what "local data" means in the first place.

One-way functions as non-split epis; preimage search as section non-existence

In categorical terms, a one-way function behaves like an epimorphism e: X -> Y that does not split: there is no s: Y -> X with e ∘ s = id_Y inside T. For hashes such as SHA-256 modeled as h: {0,1}^* -> {0,1}^256, preimage search asks for a section on the fiber over a given y. Classically, the claim "SHA-256 is preimage-resistant" can be read as:

1. h is an epi (surjective onto its image).
2. There is no computable section s in T (non-split epi).
3. Even after passing to refinements of T that include more classical resources (RAM, subexponential circuits), e remains non-split.

Grover's algorithm shows that in a *quantum* topos T_q (or, more modestly, in a fibred category that adds amplitude dynamics), the effective height of the search problem collapses from O(N) to O(sqrt(N)): a *path-shortening* homotopy in the witness space. Yet Grover alone still assumes an oracle and does not produce an unconditional section for real-world hashes; it "lowers the

hill" without changing its topology. The deeper threat comes when quantum structure supplies physically checkable witnesses that *retype* the morphisms—turning some formerly non-split epis into split epis within T_q, or annihilating the relevant obstruction classes.

Why verifiable quantum tasks change the epistemic status of outputs (Google Research)

Historically, "quantum advantage" raised a trust problem: even if a device outputs something classically intractable, how do we *know* it did the right thing? Google's *Quantum Echoes* addresses this by implementing an echo protocol for out-of-time-order correlators (OTOCs): run U, insert a controlled perturbation, then apply U^{-1}; a correct computation yields a measurable *return* that classical simulators cannot reproduce at comparable cost. The key leap is verifiability: a *physical* test supplies a witness for correctness—an observable "echo"—that external parties can check without simulating the entire process. Experiments on Willow report a ~13,000x speedup over state-of-the-art classical methods for this validated task class and motivate near-term applications in Hamiltonian learning and NMR, where verified quantum data become inputs to downstream pipelines (e.g., molecular modeling).

Categorically, echo-verification behaves like constructing a zigzag X <-U^{-1}- X -P-> X -U-> X whose composite yields a witnessable invariant. That zigzag is a *section-like* artifact in the measurement sheaf: it acts as a *computational certificate* that did not exist classically at feasible cost. Once the evidence sheaf changes, cohomology can change: obstructions that blocked global sections may vanish because the notion of "local checkability" is stronger. In short:

- Classical hardness: "no global section; obstructions persist under classical refinements."
- Echo-verified quantum tasks: "new, physically testable local data; some obstructions may now glue."

This does not mean SHA-256 is broken today. It means that shape—not speed—sets the frontier: as verifiable quantum protocols expand, the *type* of evidence available to adversaries changes. Post-quantum design should therefore target *topos-invariance*: primitives whose security class remains intact under echo-style liftings, error-corrected evolutions, and other structure-altering functors—exactly the axis Willow's results illuminate.

Category Theory of One-Way Functions

1) Faithful vs. full functors: compression morphisms and loss of invertible data

Let C be a category of *structured* data (strings with timing/memory traces, or bitstrings with provenance) and computable morphisms. Let D be a "coarser" category that forgets some of that structure (just bitstrings), and let

U : C -> D
be the corresponding forgetful functor.

- Faithful means U reflects equality of arrows: if U(f)=U(g) then f=g. Intuitively, faithful forgetfulness keeps enough information about *maps between objects* that distinct algorithms remain distinct after forgetting structure.
- Full means U is surjective on hom-sets: for any two objects X,Y, every D-arrow U(X)->U(Y) *lifts* to some C-arrow X->Y. Fullness is much stronger: it says "no morphisms are lost by forgetting structure."

In cryptography the relevant forgettings (from "rich process state" to "bare output") are faithful but not full. They preserve distinguishability of implemented algorithms (we can still tell SHA-256 from SHA-512), yet they *lose* many potential inverses: a D-arrow "candidate inverse" often has no lift to C because the requisite preimage-selecting computation cannot be realized under the resource and structure constraints of C.

A hash h : X -> Y can be factored as:
X --(q)--> Q --(π)--> Y
where q is a *quotient/compression* (coequalizer of an equivalence ~ on X), and π is an encoding into the fixed output space Y. At the functorial level, this looks like applying a left adjoint (a reflector/quotient) that *identifies* many inputs. The passage through q loses the data needed to reconstruct inverses; U being faithful-not-full explains why *some* right-inverses exist abstractly in D but do not lift to computable right-inverses in C.

Punchline: One-wayness is the categorical shadow of "faithful-not-full forgetting composed with a quotient." Compression is not just high cost to invert; it is *non-liftability of right inverses* across U and the quotient.

---

2) SHA-256 as a collapse morphism (entropy reduction) and fibrational boundaries

Model SHA-256 as a regular epimorphism (coequalizer)
h : {0,1}^* -> {0,1}^256.
It induces a partition of X = {0,1}^* into fibers X_y = h^{-1}(y). Categorically, h collapses vast subobjects to a point in Y. Information-theoretically, this is extreme *entropy reduction per output symbol*: many microstates (messages) map to a single macrostate (digest).

To track *feasible* witnesses (preimages we can actually find), use a fibration
p : E -> Y,
where the fiber E_y encodes *admissible search states/algorithms* for the preimage of y under

the resource model (time, memory, queries, side information). Think of p as the "witness fibration."

- The fibrational boundary of h is the region where fibers go from "algorithmically thin" (structured families we can traverse) to "algorithmically thick" (astronomical, pseudo-randomly distributed, no low-cost structure to exploit). Classical security statements assert that for random-looking h, almost all fibers lie beyond that boundary.
- Differential, algebraic, or side-channel structures (length extension, Merkle–Damgård quirks, leakage) *deform* the fibration p, thinning some fibers and creating liftable paths from Y back into X. Grover-style amplitude amplification "shortens" paths but does not—by itself—change the quotient nature of h; the collapse morphism remains, only its *effective height* drops from N to sqrt(N) queries.

Thus, SHA-256 as a collapse morphism means: (i) regular-epi identification of huge preimage classes; (ii) an induced witness fibration whose generic fibers are so thick that no cartesian lift (no section or even local cleavage with feasible transport) exists under classical constraints.

---

3) Obstruction to global sections as the categorical face of "hardness"

Let F be a presheaf (or sheaf) over Y assigning to each open "neighborhood of outputs" $U \subseteq Y$ the *set of locally valid witness procedures* for preimages in U (e.g., meet-in-the-middle strategies valid on a tiny subset, or collision gadgets under special padding). Local sections of F are *partial solvers*. "Breaking" preimage resistance amounts to producing a global section s : 1 -> Γ(F) that *uniformly* selects a (feasible) witness for every y.

- In Cech/derived terms, the failure of local sections to glue yields an obstruction class [o] in H^1(Y, F) (or a suitable cohomology built from the computational site). When $[o] \neq 0$, local tactics cannot be coherently assembled into a single, globally computable inverse/witness selector.
- Hardness = persistence of [o] under all *admissible refinements* of the classical site: larger memory covers, deeper time covers, more sophisticated heuristics. As long as the evidence presheaf does not gain new glueing data, the obstruction remains non-vanishing.
- A collapse of hardness corresponds to the annihilation of [o]: either because the site is enriched (e.g., oracle access, side information, leakage), or because the *notion of local verifiability changes* (e.g., physically checkable witnesses). In quantum settings with *verifiable* subroutines (echo-type certificates, fault-tolerant invariants), the presheaf F is *strengthened*: what counts as a local witness expands to include artifacts that were

infeasible to certify classically. This can convert previously incompatible local sections into compatible ones—killing the obstruction class and opening a (cartesian) lift along p.

Bottom line:

- One-wayness = non-split epi (no section).
- Cost-based security = “no section at feasible height.”
- Categorical/homological hardness = “no *global* section because the obstruction class does not vanish.”
- Structural threats (including verifiable quantum subroutines) are dangerous not merely because they reduce height (work), but because they change the sheaf and the site, potentially making [o] = 0—i.e., turning *shape* against the scheme rather than pushing harder on *speed*.

HoTT of the Preimage Problem

1) Dependent types for instances–witnesses; uninhabited preimage types classically

Let X be the type of messages, Y the type of digests, and h : X -> Y a hash.
For a fixed output y : Y, define the *preimage dependent type*:

Preim(y) := Sigma (x : X). Id_Y(h x, y)

- An element of Preim(y) is a pair (x, p) with x : X and a path p : h x = y. In standard HoTT, Id encodes equality (path) in Y.
- *Computable preimage selection* is a (partial) section s : Y -> X together with a proof that h (s y) = y, i.e., a term in:

Sect := Pi (y : Y). Sigma (x : X). Id(h x, y)

- Classical “preimage resistance” can be expressed as: for “most” y, the type Sect has no term *in the feasible subuniverse* (see modality below). Equivalently: for generic y, Preim(y) may be inhabited in principle (there exist many x) but *no inhabitant is computably recoverable* under the given resource discipline.

To track feasibility, equip the theory with a *modality* M (a reflective subuniverse) that interprets “physically/algorithmically realizable.” Think of M as sending a type A to its “feasible reflection” M A with unit eta_A : A -> M A. Examples:

- M_cl = “classically computable under cost model C”

- M_q = “quantum-realizable under cost model Q”

Then *classical uninhabitance of witnesses* reads:

not (M_cl Preim(y) is inhabited) for generic y

even though (Preim(y) is inhabited) propositionally (there exist many x such that h x = y).

Thus, the instances–witnesses story is cleanly dependent-typed: an *instance* is y : Y; a *witness* is (x,p) : Preim(y). Resistance is the systematic failure to produce M_cl-terms of Preim(y).

---

2) “Constructive existence” vs. mere existence: when a witness is physically realizable

HoTT distinguishes between *mere existence* (propositionally truncated) and *constructive existence* (an explicit term):

- Mere existence of a preimage for y is captured by the propositional truncation:

||Preim(y)|| : the (-1)-truncation of Preim(y).

This forgets the actual witness and retains only the fact “there exists at least one.” Hashes are designed so that for each y, ||Preim(y)|| is trivially inhabited (many preimages), but *this is epistemically useless for inversion*.

- Constructive existence is an inhabitant of Preim(y) in the *feasible* world:

w_y : M Preim(y)

The computational question is precisely whether we can *lift* from mere existence to constructive existence *inside M*:

exists_to_constructive_M : ||Preim(y)|| -> M Preim(y) ?

Security asserts that no such natural transformation exists for generic y when M = M_cl. In categorical terms, the unit/truncation composites cannot be split by any feasible section.

Two practical refinements:

1. *Search-with-certificate types.* Often we accept a witness only with a verifier. Package a verifier v : (x : X) -> Dec(Id(h x, y)) and ask for:

Wit_ver(y) := Sigma (x : X). v x = inl p

Constructive existence then means an element of M Wit_ver(y). Classical soundness is preserved, but realizability may change if we move from M_cl to M_q.

2. *Modal naturality.* If there exists nat : (y : Y) -> M Preim(y) that is *natural* in y (respects reindexing along maps on Y), then we effectively have a computational section in the M-world. Cryptographic hardness says such nat does not exist for M_cl; the quantum question is whether certain M_q-augmented tasks (e.g., echo-verified subroutines) permit partial nat on non-negligible families of y.

---

3) Amplitude amplification as curvature on a type space (intuition for Grover)

Grover search rotates amplitude in a 2D subspace spanned by |good> and |bad>, achieving O(sqrt(N)) queries. In a HoTT-flavored intuition, think of *search* as *transport* across a "witness bundle" over Y:

- Consider the family Preim(y) over y : Y as a fibration. A search procedure induces a path in a *control space* C (queries, oracles, reflections) that attempts to land in a *contractible* subspace of Preim(y) (a locator for a specific x with a checkable path p : h x = y).
- *Amplitude amplification as curvature.* Endow the control space with a connection whose parallel transport acts on "belief amplitude" about membership in Preim(y). Flat transport (no curvature) corresponds to unbiased random walks: expected hitting time ~ N. Grover's reflections impose a structured, periodic *curvature* that bends the transport toward the good fiber, shortening the geodesic to O(sqrt(N)). Formally, one can model the two-reflection operator as a discrete geodesic flow on the projective space of states; the sectional curvature induced by the iterate makes the "shortest path" to the contractible target subspace quadratically shorter than in the flat metric.
- In dependent-type terms, define a *success predicate* Good(x,y) := (Id(h x, y))^. (possibly truncated/decidable). The algorithm implements an endomap G on the space of distributions over X such that the *angle* to the subobject classified by Good decreases at a controlled rate:

theta_{t+1} = theta_t - 2*alpha, with alpha set by oracle reflections,

so that after O(1/alpha) steps we enter a neighborhood where Good holds with high probability and a verifier yields (x,p) : Preim(y).

Limits and structure:

- Grover changes *metric height* but not the *type-theoretic topology*: Preim(y) did not become easier to prove inhabited; we only reduced the cost to *discover* an inhabitant when an oracle for Good is available.

- Structural (echo-verified) quantum subroutines can *enrich the modal world M_q itself*: they may supply physically checkable local data that alter the *sheaf of witnesses* over Y. That can, in principle, change glueing behavior (obstructions) rather than only shortening geodesics.

Takeaway.

- HoTT cleanly separates “there exists” (||Preim(y)||) from “here is” (Preim(y)).
- Security claims live in a *modal* subuniverse M encoding feasibility; classical hardness is the failure to realize Preim(y) in M_cl.
- Grover provides a geometric intuition: curvature that shortens paths to witnesses without changing the underlying dependent type. Quantum protocols that also change what counts as *verifiable local evidence* potentially modify M itself—opening a qualitatively different route from cost reduction to *shape change*.

Grover’s Algorithm as Functorial Lifting

1) Grover as a diagonal on subobjects; square-root law as height collapse

Let h : X -> Y be a hash and fix y : Y. The preimage predicate

Good(x) := [h(x) = y]

picks out a subobject G ↪ X. In Set, G is classified by the characteristic map χ_G : X -> Ω (truth values). Grover’s move is a functorial lifting of this subobject along a semantics functor

F : (Set, subobjects) -> (Hilb, projectors)

that sends predicates to orthogonal projectors: χ_G |-> Π_G. With black-box access to χ_G (an oracle that marks “good” points), we realize two reflections in Hilb:

R_G = I - 2 Π_G (reflect about the good subspace)

R_ψ = 2 |ψ><ψ| - I (reflect about the initial state |ψ>)

Their product Q = R_ψ R_G is a rotation in the 2D plane spanned by {|good>, |bad>}. After k applications,

sin(θ_k) = sin((2k+1)θ_0)

so amplitude concentrates on G in k ~ O(sqrt(N/M)) steps (N = |X|, M = |G|). In query-complexity terms, this is a height collapse from N/M to sqrt(N/M): the “vertical” cost of

climbing to the witness subspace drops by a square root once the subobject has been lifted to a projector.

Interpretation. The diagonal/characteristic of a subobject in Set becomes a physical *involution* in Hilb. Grover's iterate composes these involutions to create a shortest-angle path to the good fiber—precisely the square-root law.

---

2) From search cost to path deformation — homotopies that shorten certificate paths

View witnesses as dependent pairs (x, p) inhabiting Preim(y) = Σ x:X . Id(h(x),y). A *certificate path* is any computable procedure that transports an initial belief state to such a witness with a verifiable equality proof p.

- Homotopic picture. The iterate Q defines a discrete flow on complex projective space P(H). In the classical (random) walk, the path to G is long and meanders (large "geodesic height"). With Grover's reflections, we *deform the path* by composing two symmetries so that each step is a deterministic rotation toward G. This is a homotopy of algorithms: same endpoints (start distribution -> witness neighborhood), but a new path whose length scales as sqrt(N/M) rather than N/M.
- Certificate locality. Because membership in G is decidable by evaluating h(x) and comparing to y, once amplitude lands in G we extract an x and a classical p : h(x)=y cheaply. Grover does not create new witnesses; it bends the search geometry so the certificate path is shorter.
- Sheaf intuition. Over output space Y, the family of preimage sets forms a bundle. Grover equips the control space (queries/oracle calls) with a curvature that biases parallel transport toward sections landing in the good fiber, shrinking the "lift distance" to a witness without altering the fiber's combinatorial type.

---

3) Limits: oracle assumptions; why Grover alone doesn't trivialize hashes

Grover is powerful but *structurally conservative*. It preserves the underlying "shape" of the preimage problem and therefore does not trivialize cryptographic resistance:

1. Oracle model. The quadratic speedup assumes an oracle that flags Good(x). For real hashes, the only honest oracle is "compute h(x) and compare to y," so each query costs a full SHA-256 evaluation. Result: preimage search drops from O(2^n) to O(2^{n/2}) evaluations—still exponential.

2. Lower bounds. Bennett–Bernstein–Brassard–Vazirani (BBBV) show the Ω(sqrt(N/M)) lower bound for unstructured search. Absent additional structure (e.g., algebraic regularities, leakage), Grover is optimal. Thus no black-box quantum algorithm can do substantially better on generic preimage problems.

3. No retyping of the morphism. Grover changes *metric height* (work) but not the *topology* of the map h. The epimorphism stays non-split: there is no functorially natural section s : Y -> X produced. In categorical terms, Grover is a lifting that acts on *predicates* (subobjects -> projectors) rather than on the quotient nature of h itself. It shortens geodesics; it does not eliminate the coequalizer that collapses many inputs to one digest.

4. Multi-target and structure do not save you (generically). If M = |G| is large (many valid preimages), cost is ~ sqrt(N/M), but for cryptographic hashes M is designed to be ~1 for random y. Known speedups (quantum walks, claw finding, collision search) alter constants/exponents for specific tasks (e.g., collisions ~ 2^{n/3}), yet preimage remains at ~ 2^{n/2} without exploitable structure.

5. Verification vs. inversion. Echo-verified quantum subroutines (e.g., Willow's Quantum Echoes) change the *epistemic* status of certain quantum outputs (you can *trust* them). But they do not by themselves yield an inverse for h. Only when quantum structure supplies new, physically checkable *local evidence* that rewires the witness sheaf (beyond black-box search) do obstruction classes risk vanishing. Grover, by design, does not perform such a retyping.

Bottom line. Grover is a clean functorial lifting of subobjects to projectors that yields a square-root height collapse via geometric path deformation. It is the right geometric lens on black-box search—but, standing alone, it does not trivialize cryptographic hashes, because it does not alter the quotient structure, the non-split epimorphism, nor the cohomological obstructions that constitute cryptographic "hardness."

Willow and Quantum Echoes: Echo-Verifiable Witness Construction

The Echo protocol: run U, perturb, run U^{-1}; compare to measure OTOCs

Quantum Echoes implements a forward–perturb–reverse loop:

1. Forward: prepare $|\psi_o\rangle$ and evolve under a random (or task-specific) circuit U to $|\psi_1\rangle = U|\psi_o\rangle$.

2. Perturb: apply a small, localized operation P (e.g., a single-qubit Pauli) to produce $P|\psi_1\rangle$.

3. Reverse: evolve back with $U^\dagger$ to obtain $|\psi_2\rangle = U^\dagger P U |\psi_o\rangle$.
4. Compare: measure an observable on $|\psi_2\rangle$ and compare to the unperturbed return $U^\dagger U|\psi_o\rangle = |\psi_o\rangle$. The contrast yields an out-of-time-order correlator (OTOC), here of the OTOC(2) family.

Operationally, the Echo return encodes the commutator growth of local perturbations under U and makes chaos/scrambling *measurable*. Willow executes this loop across 100+ qubits with high circuit depth, so that classical simulators struggle, yet the *presence/absence/size* of the return is captured by a simple, robust measurement at the end.

Why echoing enables verifiability (Loschmidt-style return as a physical section)

Classical "quantum advantage" results often lacked an external *certificate*; verification required re-simulation. Echoes supply a Loschmidt-style return: if hardware actually implements U, then $U^\dagger$ cancels it, so deviations after the perturbation are *physically checkable evidence* that a specific non-classical evolution occurred. In categorical terms, the echo measurement behaves like a *section-like witness* in a measurement sheaf: a short, observable test that glues local behavior (U, P, $U^\dagger$) into a globally checkable invariant—no large classical compute needed. That is why Google frames the result as verifiable quantum advantage on hardware rather than mere speedup.

Empirics: ≈13,000× speedup for OTOC(2)-like tasks; Nature publication claims

Google reports that Willow, running Quantum Echoes, achieved about a 13,000× performance advantage over best-in-class classical approaches on this validated OTOC(2) task class, with experiments spanning ~103 logical participants (qubits under evolution/echo). The team positions this as the first verifiable quantum advantage *on hardware*, with a companion Nature paper detailing circuits, benchmarks, and verification. Independent coverage (news/industry press) repeats the 13,000× figure and highlights that the advantage is tied to *verifiable physics workloads* rather than synthetic benchmarks.

Practical hooks: Hamiltonian learning/NMR and molecular structure inference

Because Echoes measures how local perturbations spread, it directly informs Hamiltonian learning and NMR-like spectroscopy: the echo signal maps couplings and correlation structure, which are the raw ingredients for reconstructing molecular graphs and material interactions. Coverage emphasizes these as the first near-term utility targets—drug discovery, materials science, and broader physical modeling—precisely because *echo-verified* outputs can be trusted as inputs to classical/AI pipelines without full re-simulation. Expect early pipelines where Willow-class devices generate validated correlators that downstream algorithms invert into couplings, constraints, or candidate structures.

One-paragraph synthesis

Echo-verifiable witness construction reframes advantage: the U → P → U† loop yields a physical certificate (the Loschmidt return) that upgrades "we ran a hard circuit" into "we produced *checkable* evidence of non-classical dynamics." Empirically, Willow's Quantum Echoes posts a ~13,000× speedup on OTOC(2)-style problems and anchors claims in a Nature paper. Practically, echo-validated correlators plug into Hamiltonian learning/NMR workflows to infer molecular and material structure—precisely the kind of shape-of-physics data that classical pipelines can trust and reuse.

Echoes as Zigzag Diagrams in the Quantum Topos

The $F \circ P \circ F^{-1}$ loop as a zigzag witnessing equivalence classes

Let Obj($\mathcal{Q}$) be objects (state-preparation contexts, code spaces, or "instances") and Mor($\mathcal{Q}$) be physically realizable evolutions. Write

- $F : X \to X$ for the forward evolution (unitary circuit U acting on the prepared object X),
- $P : X \to X$ for a localized perturbation (small, sparse operator),
- $F^{-1} : X \to X$ for the time-reversal (U^{-1}).

The Echo protocol composes these into $E = F^{-1} \circ P \circ F$. Rather than viewing E as a single arrow, read it as a zigzag in the slice category $\mathcal{Q}/X$:

X --F--> X --P--> X --F^{-1}--> X

A zigzag class [F ⇄ P] is the homotopy-invariant content of "conjugating a local act by a global flow." Echo measurements compare the endpoints of this zigzag with the identity arrow id_X. When the return is high, $E \simeq id_X$ in the homotopy category Ho($\mathcal{Q}$); when scrambling occurs, E moves into a distinct equivalence class detectable by the chosen observable. In this sense, echoes witness equivalence classes: the protocol constructs, then *certifies*, whether P is trivialized (up to homotopy) by conjugation through F.

Categorically, you can regard the conjugation functor

Ad_F : End(X) -> End(X), Ad_F(P) = $F^{-1} \circ P \circ F$

and the echo readout as a natural transformation comparing Ad_F to the identity functor on End(X). A nontrivial echo indicates that Ad_F carries local generators into a measurably different component; a strong return indicates that Ad_F fixes P up to a contractible neighborhood. Thus,

the U → P → U^{-1} loop is a *computational* zigzag whose measured closure (how near it lands to id_X) *proves* which homotopy class we occupied.

Noise as a sheaf deformation; logical qubits as near-global sections under refinement

Let $\mathcal{U}$ be a cover of the experiment by contexts of control and measurement (patches of calibration, gate neighborhoods, readout frames). Associate to each patch U ∈ $\mathcal{U}$ a set (or space) of *locally admissible behaviors*—circuits that meet fidelity specs, perturbations we can implement, and observables we can reliably read. This assignment forms a presheaf $\mathfrak{E}$ on the site of experimental contexts:

$\mathfrak{E}$(U) = {locally valid (F,P,F^{-1}, observable) tuples over U}

Noise and crosstalk deform the restriction maps of $\mathfrak{E}$. In a noiseless ideal, local tuples glue to a global tuple whose echo closes (E ≃ id_X) or deviates in a quantifiable way; the presheaf is close to a sheaf. Under noise, compatibility conditions fail: sections that look valid on overlapping patches disagree when glued. Mathematically, this is a sheaf deformation: the same base site, but altered restriction maps and gluing axioms due to stochastic channels and coherent errors.

Error correction introduces a *refinement* of the site. Replace raw physical patches by code patches (stabilizer neighborhoods, decoding regions). Pull $\mathfrak{E}$ back along this refinement and push forward through decoding; the new presheaf $\mathfrak{E}$_log has improved gluing because syndromes *transport local inconsistency into known defect channels*. In favorable regimes, logical qubits appear as (near-)global sections:

- A logical state $|\bar{\psi}\rangle$ corresponds to a section s ∈ Γ($\mathfrak{E}$_log) that is stable under restriction and re-gluing up to correctable deviations.
- The echo zigzag at the logical layer, Ad_$\bar{F}$($\bar{P}$), remains in the same homotopy class across the cover because decoding re-identifies small local mismatches before they obstruct gluing.

From a fibration viewpoint, let p : $\mathfrak{W}$ → $\mathcal{Y}$ be the witness fibration, where the base $\mathcal{Y}$ indexes measurement contexts and the fiber collects feasible “echo-witnesses” (return statistics, OTOC estimators, certificates). Noise perturbs p by thickening fibers and bending transport (parallel arrows no longer compose exactly). Error correction plus echo re-straightens transport: decoding defines cartesian lifts that map noisy local witnesses back into a consistent logical witness, shrinking the coboundaries that once prevented global sections. Practically, this is why echo protocols pair so naturally with fault tolerance: the Loschmidt-style return supplies a section-like certificate, and code-based refinement ensures those certificates glue—turning fragile local evidence into a coherent, system-wide statement.

Summary.

- The U → P → U^{-1} echo is a zigzag whose measured closure distinguishes homotopy classes of local perturbations under global dynamics.
- Noise deforms sheaves of valid operations/measurements, breaking gluing; error correction refines the site so that logical behavior re-acquires (near-)global sections.
- In this refined, "quantum topos" view, echoes are not just signals—they are witness constructors whose ability to glue across patches is the operational meaning of verifiability.

Collapse Conditions for Cryptographic Types

When obstruction classes vanish: from hard preimage types to contractible ones

Let $h: X \to Y$ be your hash and for each $y: Y$ the dependent preimage type

$$\mathrm{Preim}(y) \;=\; \Sigma(x: X).\, \mathrm{Id}(h\,x, y).$$

"Hardness" is the non-existence of a *feasible* global section $s: Y \to X$ together with proofs $p_y: \mathrm{Id}(h(s\,y), y)$, formalized as the persistence of an obstruction class $[o] \neq 0$ in a cohomology built from the evidence presheaf $\mathcal{F}$(local witness procedures) over a classical site S_{cl}.

Obstruction classes *vanish*—and preimage types become *effectively contractible*—when any of the following shape changes occur:

1. Modal enrichment (feasibility widens). Move from the classical feasibility modality $M_{\mathrm{cl}}$to a richer one M_{qv}(quantum + verifiable). If there exists a natural transformation

$$\eta: \;\; ||\; \mathrm{Preim}(-) \;|| \;\;\Longrightarrow\;\; M_{\mathrm{qv}}\, \mathrm{Preim}(-)$$

on a non-negligible family in Y, then the truncated "there exists" lifts to constructive, *physically realizable* witnesses. Gluing succeeds; $[o] = 0$.

2. Sheaf strengthening (more local evidence). Echo-verifiable subroutines add short *physical* certificates to $\mathcal{F}$. On a good cover $\mathcal{U}$of output space, the Čech coboundary groups $\check{H}^1(\mathcal{U}, \mathcal{F})$ collapse (e.g., $\mathcal{F}$ becomes flabby/acyclic), so local sections now glue globally.
3. Exactness under base change (descent holds). If a covering functor $J: \mathsf{S}_{\mathrm{qv}} \to \mathsf{S}_{\mathrm{cl}}$ induced by verification satisfies Beck–Chevalley/exactness, then global sections computed

upstairs (quantum-verified world) *descend* to the classical base. Practically: a logically protected, echo-certified witness at the code level descends to a classical string $x$with a checkable p, killing the obstruction.

4. Contractible retyping of fibers. Added structure (invariants, symmetries, side information) makes each relevant $\mathrm{Preim}(y)$ equivalent to a pointed, path-connected type with unique witness up to homotopy. Once the fiber is contractible, giving *any* point produces a canonical section.

Net effect: the failure to glue (hardness) is a *shape fact*. When the measurement/witness sheaf is strengthened or the site is refined, the shape changes; hardness can collapse even if raw work factors haven't.

Distinguish "break by qubits" from "retype by structure": echoes vs. oracles

- Break by qubits (height reduction). Grover/BBBV lives here. You keep the same site S_{cl} and the same presheaf $\mathcal{F}$, but use more "vertical power." Costs drop from $2^n$to $2^{n/2}$ for preimage, yet the epimorphism h remains non-split. No new sections; only shorter geodesics to occasional witnesses.
- Retype by structure (shape change). Echo protocols, logical encoding/decoding, and other *verifiable* quantum routines alter the *typing* of evidence: they add physically checkable local data and functorial comparisons (e.g., the Loschmidt return) that weren't available classically. That is a *retyping* of the witness presheaf and/or a refinement of the site. Obstructions can vanish; sections can now exist *as objects of the new world* and sometimes descend.

Think of it as oracles vs. echoes. Oracles grant speed in a fixed space; echoes change the *space of certificates* you're allowed to use.

Adversarial adjunctions that turn non-sections into sections (threat model upgrade)

An attacker can *functorially* upgrade their world so that right inverses appear canonically.

- Adjoint completion. Suppose there is an adjunction

$$A \dashv U\colon \mathsf{C} \rightleftarrows \mathsf{Q},$$

with U "forgetful to classical," and A "add structure" (encode, echo-verify, enrich witnesses). A classical hash $h$lives in C. If Asends $h$to $A(h)$ whose preimage fibration admits a section (e.g., due to echo-certified invariants), then the unit η: Id $\Rightarrow U\ A$ yields a *canonical* map

$$s\colon Y \xrightarrow{\eta} UA(Y) \xrightarrow{U(\tilde{s})} UA(X)$$

that composes with $h$to the identity *in UA-world*, i.e., a section after enrichment, which may descend.

- Monadic capture (Beck). If U is monadic, then algebras for the monad $T = UA$ encode "programs with verification privileges." A T-algebra structure on Ycan force existence of algebraic sections $Y \to X$ not available in plain C. Adversaries aim to *force monadicity* by coupling devices (echo/verifier pairs) that act like algebra structure.
- Descent along verified covers. Provide a cover of Y by contexts $\{U_i\}$ where each U_i admits an echo-verified local section $s_i$and verified cocycle equalities on overlaps. If those equalities live in a rigid, low-noise logical layer, the Čech 1-cocycle is trivial, so the $s_i$glue into a global s. What was a non-split epi becomes split under the verified cover.
- Leakage-as-left adjoint. Side-channels or protocol hints can be modeled as a left adjoint Lthat *adds observables*. If $L(h)$ factors through a lower-entropy quotient, the induced fibers shrink and may become contractible; a section appears *because the type changed*, not because we "searched harder."

Threat-model takeaway. Defenses that only price *height* (more bits, larger work factors) ignore adjoint attacks that change the *site* or *sheaf* of evidence. Post-quantum designs must require *topos-invariance*: security predicates preserved under (i) enrichment adjunctions (encode/verify), (ii) descent along verified covers (fault-tolerant logical layers), and (iii) leakage-adding functors. If your primitive's hardness is merely "no section in $\mathrm{S_{cl}}$," an adversary may move you to $\mathrm{S_{qv}}$ where $[o]$ vanishes and a section *exists by construction*.

Post-Quantum Security as Topos-Invariance

Protocols as objects in a fibred category; security as preservation under retyping

Model a cryptographic protocol as an object Pliving over a base category of contexts (keys, randomness, transcripts), with a witness fibration $p\colon \mathcal{W} \to \mathcal{C}$ whose fibers encode feasible adversarial procedures. *Topos-invariant* security means the predicate "no feasible global section exists" (no uniform inversion/extraction) is preserved under retypings—i.e., along structure-adding functors that change what counts as local, checkable evidence (e.g., echo-verified

quantum subroutines). In this view, Willow's Quantum Echoes enrich the site with physically checkable certificates, but a post-quantum-sound Pkeeps its "no-section" property under such liftings.

Sheaf-theoretic views of commitment, randomness, and transcript consistency

Treat per-round artifacts (commitments, coins, messages) as a sheaf $\mathcal{E}$over the cover of operational contexts (devices, timing, faults).

- Commitment soundness: local binding/hiding data must glue across overlaps; attacks manifest as non-trivial 1-cocycles (inconsistencies that still pass local checks).
- Randomness: model as a sheaf of extractable entropy; leakage/side-info deforms restriction maps.
- Transcript consistency: verifiers check that local constraints assemble into a global section; echo-style physical checks add new local sections that were absent classically. Designing PQC so that these added sections do not trivialize the obstruction to inversion is exactly the "invariance under retyping" requirement. Echoes change *evidence geometry*; sound protocols remain hard despite that change.

Design goal: primitives whose hardness is stable under echo-like liftings

Concretely, aim for primitives where hardness proofs are robust to:

1. Certificate enrichment: adversaries may obtain short, *physically verifiable* correlators (echo returns, OTOCs) that can be validated without classical re-simulation. Security should remain non-split even when these certificates are available.
2. Descent from logical layers: if an attacker computes in an error-corrected, verified layer and descends results to the classical transcript, obstruction classes should persist (no global witness becomes available after descent).
3. Exactness under base change: proofs withstand base changes from the classical site to an echo-verified quantum site (no spurious sections appear because the sheaf became "too flabby").
   Practically, that means prioritizing PQC families (KEMs/DSAs) whose security reductions don't rely on structural assumptions likely to be weakened by physically verified correlators—exactly the thrust of NIST's standardization (ML-KEM/Kyber, ML-DSA/Dilithium, SLH-DSA/SPHINCS+, plus HQC).

Standards context & near-term reality

Two guardrails for readers:

- What Willow does—and doesn't—imply. Google's team presents *verifiable* advantage for OTOC-style tasks with about 13,000× speedup over top classical methods and a Nature paper backing the claim. This upgrades trust in certain quantum outputs (useful for Hamiltonian learning / NMR), but does not equate to breaking mainstream public-key crypto.
- PQC is prudent now. NIST finalized the first PQC FIPS (203/204/205) in Aug 2024 and added HQC in Mar 2025; migration guidance emphasizes "crypto-agility" given "harvest-now-decrypt-later" risk. Industry and standards bodies recommend inventorying crypto, prioritizing ML-KEM/ML-DSA/SLH-DSA adoption, and planning staged cutovers.

Bottom line: Treat security as shape, not speed. A post-quantum scheme should keep its "no global section" property even after the evidence sheaf is enriched by echo-verifiable certificates and logical-layer descent. Willow strengthens the case for PQC—not because it breaks RSA today (it doesn't), but because it shows how *retypings* of evidence will arrive before raw qubit counts are sufficient to mount large-scale cryptanalytic attacks.

From Resonance to Echoes: Verification as Epistemic Upgrade

Why verifiable advantage matters more than raw advantage for downstream science

"Raw" quantum advantage has usually meant *hard to simulate*, which still leaves labs and industry asking: can we trust the numbers enough to build on them? Echo-verified workloads flip that axis. By running U, inserting a controlled perturbation P, and reversing with $U^{\dagger}$, Willow creates a *Loschmidt-style return* that serves as a short, physical certificate that the intended non-classical dynamics occurred—no giant classical re-simulation required. That makes the result auditable, not just fast. Google explicitly frames Quantum Echoes as the *first verifiable quantum advantage on hardware*, elevating outputs from "impressive but opaque" to "actionable scientific evidence." This is the epistemic upgrade: verification compresses the trust gap, so quantum data can become inputs to real pipelines rather than one-off demos.

Data you can trust: enabling virtuous loops in AI, chemistry, and materials

Once you can certify correlators (e.g., OTOC(2)) at scale, you unlock tight feedback loops: (1) quantum hardware produces echo-verified signals; (2) classical/AI models ingest them to infer Hamiltonians, couplings, or candidate molecular geometries; (3) predictions guide the *next* targeted quantum runs. Because the certificate is physical (the echo return), each loop step

inherits credibility without full classical verification—exactly the kind of "trustable intermediate" that accelerates drug discovery, materials design, and broader physics modeling. Google highlights these hooks—Hamiltonian learning, NMR-like spectroscopy, and molecular structure inference—as near-term beneficiaries of Echoes, underlining that verifiability (not just speed) is what lets quantum outputs slot cleanly into downstream science and ML.

Pointer: For a succinct programmatic statement of the result and its claims (Willow, Echoes, and "verifiable quantum advantage"), see Google's announcement post.

The Cryptographic Singularity, Recast

Not a date but a limit object: as echo-style functors proliferate, hard types flatten

The "singularity" worth naming here is not a calendar event or qubit threshold; it is a limit object in the category of evidence—what you get when you close the classical world under *structure-adding functors* (encode, evolve, perturb, echo, decode) that supply physically checkable certificates. Each functor enlarges the presheaf of admissible witnesses; as these accumulate (closure under composition, products, descent from logical layers), obstruction classes that once prevented global inversion tend to vanish. In homotopy terms, families of preimage types that were formerly rugged become contractible up to verified equivalence: not because we brute-forced them, but because the *shape of admissible proof objects* changed. The "cryptographic singularity" is thus the colimit of verification power—the point where, for a broad class of target types, existence ⇒ constructive, verifiable existence under the enriched site. What collapses is not merely *cost*, but the very topology of hardness.

Predictive markers: growth of echo-verifiable workloads; cross-domain witnesses

How would we know we are asymptoting toward that limit object? Look for these operational invariants:

1. Echo density (E): fraction of quantum workloads that ship with short, *hardware-native* certificates (Loschmidt returns, stabilizer-verified invariants) usable by third parties *without re-simulation*. Rising E indicates the witness presheaf is getting "flabby" (easy to glue), a precondition for obstruction collapse.
2. Descent ratio (D): proportion of verified logical-layer results that descend to classical artifacts (bitstrings, transcripts) retaining their certificate semantics. High D means verified sections upstairs assemble into global classical sections—a direct threat vector to type-based hardness.

3. Cross-domain witness reuse (X): the rate at which a certificate produced for physics (e.g., echo-validated correlators) is repurposed to reduce hardness in non-physics domains (e.g., cryptanalysis, protocol forensics, code-equivalence). X>0 signifies we're accumulating universal witness constructors, not bespoke proofs.

4. Exactness under base change (β): frequency with which security arguments *fail to preserve* "no-section" after moving from the classical site to an echo-enriched site (base change). Increasing β is the categorical symptom of retyping outpacing parameters.

5. Certificate-to-break elasticity (ε): marginal reduction in effective hardness per additional certified observable or echo loop. If ε remains bounded away from zero across primitives, you are in a flattening regime: each new functorial certificate predictably shortens witness paths.

6. Standards backpressure (σ): how often protocol revisions are forced not by larger key sizes but by evidence-model changes (new verification channels, logical encodings, composable certificates). Growing σ marks a shift from "speed threats" to "shape threats."

Implication. Track $E, D, X, \beta, \varepsilon, \sigma$ across time. A sustained rise across several of these markers indicates the ecosystem is converging toward the limit object of verifiability, where echo-style functors—composed with error-corrected descent and classical post-processing—*systematically* turn non-split epis into split ones on meaningful subclasses. That is the cryptographic singularity recast: security fails by retyping, not by raw qubits.

Conclusion: Secure the Shape, Not the Speed

Security as geometric robustness to adversarial retypings.
Treat protocols as objects whose "hardness" is a shape property: the non-existence of feasible global sections in a given evidence topos. If adversaries can *retype the site*—adding physically checkable certificates, logical encodings, or descent mechanisms—then obstructions that once blocked inversion may vanish. Post-quantum security should therefore be framed as topos-invariance: the security predicate must persist under structure-adding functors (encode → evolve → perturb → echo → decode) and still deny a global section after such liftings. Speedups alone (qubits, query counts) are secondary; what matters is whether the shape of admissible evidence can be altered to make witnesses glue.

Willow as an early witness that verification is the decisive axis.
Google's "Quantum Echoes" result repositions advantage from "faster" to verifiable: the U → P → U^{-1} echo yields a short, *physical* certificate (Loschmidt-style return) that outsiders can check

without re-simulating the full evolution. That turns some previously local, fragile signals into globally usable evidence artifacts, a change in the evidence sheaf—not merely in compute. This is why Willow's claim is milestone-class: it demonstrates *verifiable* quantum advantage on hardware, with a ≈13,000× reported speedup for OTOC-class tasks and a peer-review anchor, while explicitly not implying modern public-key breakage today. Verification—not raw speed—opens the door for downstream science and, potentially, for future retypings of cryptographic evidence.

Program of work.

1. Topos-stable primitives. Prefer KEMs/DSAs whose reductions remain sound under echo-style certificate enrichment and logical-layer descent; prove non-existence of sections is preserved under such base changes.
2. Echo-aware threat models. Model adversaries who can generate short, physically verifiable correlators and compose them with error correction and classical post-processing; require that these do not annihilate obstruction classes.
3. Composable proofs. Build security arguments as gluing theorems over covers of operational contexts (devices, timing, noise) so that added verification channels cannot create a global section from locally consistent fragments.
4. Pragmatics now. Continue PQC migration (NIST FIPS 203/204/205; HQC selection) and crypto-agility planning; Willow's team itself notes their hardware does not break RSA today, reinforcing that PQC adoption is the correct near-term hedge while verification-centric quantum workloads mature.

Bottom line: design for shape robustness. As echo-verifiable workloads proliferate, the winning cryptography won't just outpace qubits—it will withstand retyping of the evidence landscape that Willow and its successors now make thinkable.

References

1. Mi, X.; Kechedzhi, K. “A verifiable quantum advantage.” *Google Research Blog*, Oct 22, 2025. Google Research
2. Neven, H.; Smelyanskiy, V. “Our Quantum Echoes algorithm is a big step toward real-world applications for quantum computing.” *The Keyword (blog.google)*, Oct 22, 2025. blog.google
3. Chen, Y.; Devoret, M. “Our quantum hardware: the engine for verifiable quantum advantage.” *The Keyword (blog.google)*, Oct 22, 2025. blog.google
4. Google Quantum AI and collaborators. “Observation of constructive interference at the edge of quantum ergodicity.” *Nature* (2025). Nature
5. Scientific American. “Google Measures ‘Quantum Echoes’ on Willow Quantum Computer Chip,” Oct 22, 2025. Scientific American
6. The Quantum Insider. “Google Quantum AI Shows 13,000× Speedup Over World’s Fastest Supercomputer in Physics Simulation,” Oct 22, 2025. The Quantum Insider
7. Science News. “Quantum ‘echoes’ reveal the potential of Google’s quantum computer,” Oct 22, 2025. Science News
8. Reuters. “Google says it has developed landmark quantum computing algorithm,” Oct 22, 2025. Reuters
9. The Guardian. “Google hails breakthrough as quantum computer surpasses ability of supercomputers,” Oct 22, 2025. The Guardian
10. Financial Times. “Google claims ‘quantum advantage’ that promises drug discovery breakthroughs,” Oct 22, 2025. Financial Times
11. Engadget. “Google says it made a breakthrough toward practical quantum computing,” Oct 22, 2025. Engadget
12. IEEE Spectrum. “Google’s New Quantum Algorithm May Actually Be Useful,” Oct 22, 2025. IEEE Spectrum
13. Google Research. “A verifiable quantum advantage,” overview post linking to the Nature paper, Oct 22, 2025. Google Research
14. Google Blog. “Meet Willow, our state-of-the-art quantum chip,” Dec 9, 2024. blog.google
15. Google Research. “Making quantum error correction work” (Willow below threshold), Dec 2024; with Nature companion. Google Research

16. The Verge. “Google says its breakthrough quantum chip can’t break modern cryptography,” Dec 12, 2024. The Verge

17. NIST CSRC. “Announcing Approval of Three Federal Information Processing Standards (FIPS) for Post-Quantum Cryptography (FIPS 203/204/205),” Aug 13, 2024. NIST Computer Security Resource Center

18. NIST Newsroom. “NIST releases first 3 finalized post-quantum encryption standards,” Aug 13, 2024. NIST

19. NIST CSRC. *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism (ML-KEM)*, Aug 13, 2024. NIST Computer Security Resource Center

20. NIST CSRC. “Post-Quantum Cryptography Standardization (status page including HQC selection, Mar 11, 2025),” frequently updated. NIST Computer Security Resource Center